Final Report

Title: Cooperative Airborne Inertial-SLAM for Improved Platform and Feature/Target Localisation

Contract Number: FA5209-05-P-0550

AFOSR/AOARD Reference Number: AOARD-05-4063

AFOSR/AOARD Program Manager: Tae-Woo Park, Ph.D.

Period of Performance: 30 April 2006 – 30 April 2007

Submission Date: 2007

PI: Salah Sukkarieh / ACFR, University of Sydney

Report Documentation Page

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **25 JAN 2008**	2. REPORT TYPE **Final**	3. DATES COVERED **30-04-2006 to 30-04-2007**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Cooperative Airborne Inertial-SLAM for Improved Platform and Feature/Target Localisation	**FA520906P0182**
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	5d. PROJECT NUMBER
Salah Sukkarieh	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
University of Sydney,University of Sydney, Bldg. J04,Sydney NSW 2006,Australia,NA,NA	**N/A**

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
AOARD, UNIT 45002, APO, AP, 96337-5002	**AOARD-064005**
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
The aim of this project is to develop cooperative guidance laws for platforms which employ Simultaneous Localization and Mapping (SLAM) algorithms as part of the information feedback to the guidance loop. In GPS denied environments SLAM is an essential navigation tool, as it can provide both a map of ground features together with location and attitude information for the sensor platform with respect to this map. The benefit of using the SLAM algorithm is that it can determine the accuracy of both platform and target locations, both of which improve as a function of feature/target revisitation or sharing of maps between various platforms

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**23**	

(2)   Objectives: Briefly summarize the objectives of the research effort or the statement of work.

The aim of this project is to develop cooperative guidance laws for platforms which employ Simultaneous Localisation and Mapping (SLAM) algorithms as part of the information feedback to the guidance loop. In GPS denied environments SLAM is an essential navigation tool, as it can provide both a map of ground features together with location and attitude information for the sensor platform with respect to this map. The benefit of using the SLAM algorithm is that it can determine the accuracy of both platform and target locations, both of which improve as a function of feature/target revisitation or sharing of maps between various platforms.

(3)   Status of effort: A brief statement of progress towards achieving the research objectives. (Limit this section to about 200 words or less.)

In the second year of the project we have extended the single vehicle active SLAM algorithms into the multi-vehicle domain. We began by developing a 6-DoF multi-UAV simulator with multi-vehicle SLAM algorithms where vehicles share map information with one another. This allowed us to study the interaction of information flows between the vehicles and to analyse how different sets of trajectories affected the information in each vehicle's localisation states and in the final map.
We then designed a centralised multi-vehicle active SLAM algorithm which provided trajectory commands to each of the vehicles in order to maximise map information. These control algorithms were then extended into two decentralised architectures; co-ordinated SLAM where vehicles make their own trajectory decisions based on maximising the shared map information, and co-operative SLAM where vehicles negotiate the best set of trajectories for the team.

(4)   Abstract: Briefly describe research accomplishments, their significance to the field, and their relationship to the original goals.

The development of multi-vehicle active SLAM, especially with its application to airborne vehicles is a complex problem as it touches on elements of estimation and information analysis, high-level trajectory control as well as algorithms for decentralisation of control and negotiation across multiple platforms. In reaching the overall project objectives we:
- o   Developed a 6-DoF multi-UAV simulator and demonstrated multi-vehicle SLAM where vehicles share information about the terrain to build a common map of the environment;
- o   Studied the evolution of localisation and map information between vehicles during different types of vehicle trajectories using the simulator;
- o   Developed algorithms for centralised control of multiple vehicles in order to maximise map information;
- o   Developed algorithms for two different decentralised control architectures for multiple vehicles (co-ordinated and co-operative control) in order to maximise map information;
- o   Coded the multi-vehicle control algorithms into the 6DoF UAV simulator and demonstrated multi-UAV active SLAM.

Multi-vehicle active SLAM is a challenging problem; the work presented here is novel in that it is the first time that the problems of multi-vehicle SLAM and active SLAM have been combined and studied together. This work is significant in that it provides a potential for multiple vehicles to operate in GPS-denied environments where no a-priori map information is available by providing localisation and navigation information to each vehicle. The vehicles can also co-ordinate their actions such as to build a more accurate map of the environment over which they operate. The algorithms allow the UAVs to co-ordinate their actions in such a way as to benefit the team as a whole without the need for a centralised control point.

(5) **Personnel Supported:** List the professional personnel supported by the contract and/or the personnel who participated significantly in the research effort.

Salah Sukkarieh – Principle Investigator
Mitchell Bryson – PhD Student, Inertial-SLAM development, observability analysis and co-operative SLAM path planning

(6) **Publications:** List peer-reviewed publications submitted and/or accepted during the contract period.

Accepted –
Bryson M., & Sukkarieh S., "Observability Analysis and Active Control for Airborne SLAM", accepted to IEEE Transactions on Aerospace and Electronic Systems, to appear October 2007.

Bryson M., & Sukkarieh S., "Co-operative Localisation and Mapping for Multiple UAVs in Unknown Environments", presented at IEEE Aerospace Conference, Big Sky, Montana, 3-10 March, 2007.

(7) **Interactions:** Please list:

(a) Participation/presentations at meetings, conferences, seminars, etc.

- IEEE Aerospace Conference, Big Sky, Montana, 3-10 March, 2007

(b) Describe cases where knowledge resulting from your effort is used, or will be used, in a technology application. **Not all research projects will have such cases, but please list any that have occurred.**

(8) **New:**

(a) List discoveries, inventions, or patent disclosures. (If none, report None.).

None

(b) Completed the attached **"DD Form 882, Report of Inventions and Subcontractors."**

None

(9) **Honors/Awards:** List honors and awards received during the contract period, or emanating from the AOARD-supported research project.

None

(10) **Archival Documentation:** This section should include a description of your work at a level of technical detail that you think to be appropriate. Submission of reprints/preprints often satisfies this requirement. If you have questions on how to prepare this section, please discuss this matter with your AOARD program manager.

See document attached.

(11) **Software and/or Hardware (if they are specified in the contract as part of final deliverables):** Include source code, brief installation and user guides.

None

Final Report - Cooperative Airborne Inertial-SLAM for Improved Platform and Feature/Target Localisation

Mitch Bryson and Salah Sukkarieh

I. INTRODUCTION

In applications where Unmanned Aerial Vehicles (UAVs) operate in environments where signals from the Global Positioning System (GPS) are unavailable (i.e. GPS denied environments), the vehicle is often left with terrain sensing as a method for self localisation. When an a-priori terrain map is unavailable, Simultaneous Localisation and Mapping (SLAM) is an essential navigation tool, as it can generate a map of ground features in the terrain while providing information about the vehicle's location and attitude with respect to this map. SLAM also has the potential to be distributed across multiple UAVs where each vehicle shares information it gathers about the state of the surrounding terrain with other vehicles in the network. Data sharing between vehicles helps each vehicle to improve its own localisation accuracy and in turn allows the construction of a larger and more accurate map of the terrain.

This final report summarises the work completed on the theory and practical results in the second year of the project "Cooperative Airborne Inertial-SLAM for Improved Platform and Feature/Target Localisation". In the first year report [4], the SLAM algorithm was presented for operation on a single vehicle. it was shown that the accuracy of the mapping and localisation estimates in SLAM was highly dependant on the control decisions made by the platform. It was shown through and observability analysis of the SLAM algorithm that localisation accuracy was dependant on the UAV making particular dynamic maneuvers while observing the terrain below. It was also shown that the accuracy of the constructed terrain map was dependant on the order of visitation and re-visitation and observation of features. Through this analysis, the first year report summarised algorithms for 'active SLAM' on a single vehicle, the concept of actively controlling the motions and trajectory of a UAV in order to maximise localisation and mapping accuracy while performing SLAM.

In this, the second year report, we extend the concept of active SLAM to the case of multiple UAVs where the team of vehicles controls their trajectories and terrain observations in order to maximise the accuracy of a global, shared map of the terrain and thus improve localisation accuracy for each vehicle. In Section II we outline a multi-UAV cooperative SLAM simulator which was used to examine results from the multi-vehicle SLAM algorithm and multi-vehicle active SLAM algorithms. In Section III we summarise the SLAM algorithm and consider two types of data fusion architectures for SLAM across multiple vehicles; firstly centralised SLAM where all information shared between multiple vehicles is communicated and fused at a central source, and secondly decentralised SLAM where information is shared amongst platforms without the need for a central communication node. In Section IV we outline algorithms for multi-vehicle active SLAM. We firstly demonstrate a centralised approach where a central control point assess the best trajectory for each vehicle to take and allocates this to each vehicle at regular intervals. Two decentralised control strategies are also shown; firstly 'co-ordinated' control where vehicle's make their own control decision based on common shared map information, and secondly 'co-operative' control where vehicles negotiate over the best trajectory for each vehicle to take that will benefit the team as a whole. Section V presents simulation results for both the multi-vehicle SLAM algorithm and multi-vehicle active SLAM control strategies.

II. MULTI-UAV COOPERATIVE SLAM SIMULATOR

In order to test and evaluate the concepts in multi-vehicle co-operation in SLAM, a multi-UAV simulator was developed. The simulator contains a complete 6-DoF flight model, low-level control system and on-board sensor model for each UAV. Multi-UAV SLAM and active SLAM algorithms were implemented into the simulator in order to test and verify the potential of multi-vehicle active SLAM. The simulator allowed us to study the information flow and accuracy of the SLAM estimates for several different vehicle trajectories which gave insight into the role of co-operation between the vehicles in active SLAM.

A visualisation of the simulation is shown in Figure 1. Figure 2 illustrates the main components of the simulation. The following subsection describe the elements of the simulation in more detail.

A. Environment

The environment in the simulation is set over a 3×3km area representing an unexplored terrain. True features that will be picked up by the terrain sensors on-board each vehicle are randomly dispersed over the terrain.
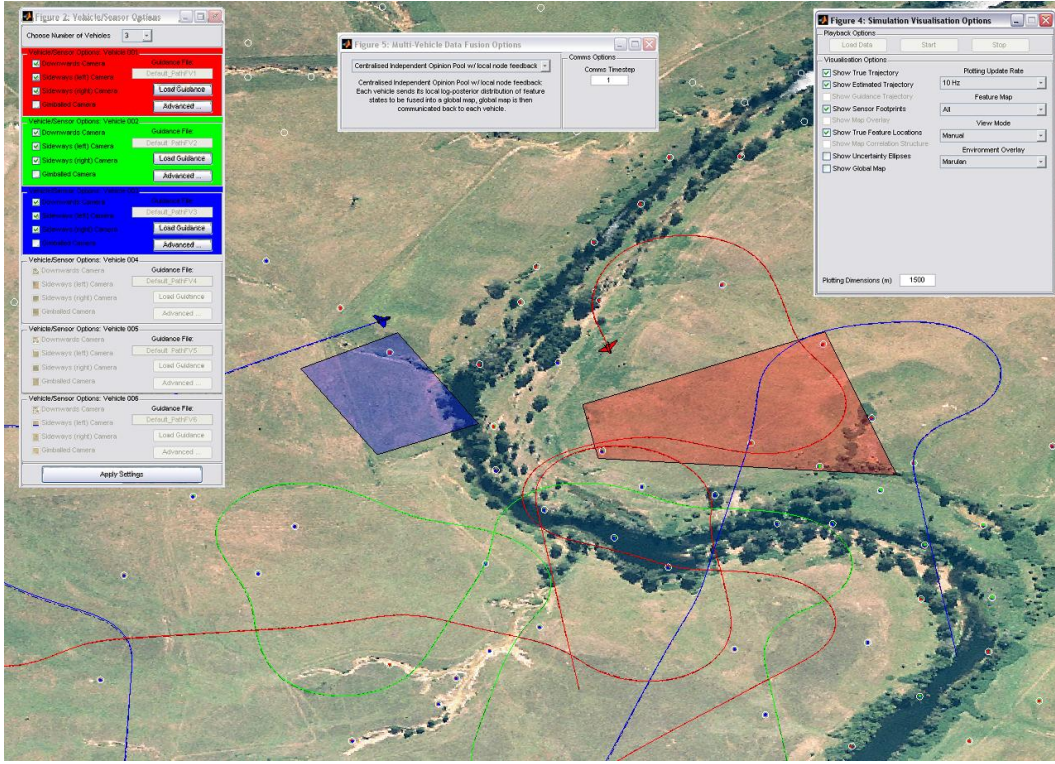
Fig. 1. Multi-UAV Co-operative SLAM Simulator Visualisation.

Vision Camera					IMU		
Sampling Rate	FOV	Maximum Range	Angular Noise	Range Noise	Sampling Rate	Accelerometer Noise	Gyro Noise
10 Hz	30^o by 22^o	400 m	1 deg (1σ)	2 m (1σ)	100 Hz	$0.05 m/s^2$ (1σ)	1 deg/s (1σ)

TABLE I

SENSOR SPECIFICATIONS: SAMPLING RATE, FIELD OF VIEW, MAXIMUM RANGE AND SENSOR ERRORS FOR THE SIMULATED IMU AND LASER-VISION TERRAIN SENSOR ON-BOARD EACH UAV.

B. 6DoF UAV Dynamic Model Simulation

Several UAVs are placed into the environment and for each UAV we simulate its motion using a 6DoF dynamic model of a real UAV, the Brumby Mk III [2], complete with aerodynamic and thrust forces on the vehicle. The dynamic model is driven by the control surface settings (throttle, rudder and elevons) and these inputs are actively controlled by the automatic guidance and autopilot systems (see next subsection).

C. Guidance and Autopilot

Low level control of each platform is broken up into two stages. The first stage is a guidance system that generates the required altitude, velocity and bank angle of the UAV to track trajectory segments that have been assigned by the central data filter. The second stage is the autopilot system which takes the desired altitude, velocity and bank angle from the guidance system and uses linear PID feedback control by moving the control surface actuators on the UAV. The guidance and autpilot modules in the simulation are derived from the actual systems on the Brumby UAV (see [2] and [3] for details).

D. Sensor Simulation

The sensor simulation generates readings for both the IMU and a laser-vision system by taking the simulation truth and adding noise. Feature observations are generated by determining which features in the terrain lie within the field of view and maximum range of the sensor. Three sensors are used, one pointing downwards and two pointing in both sideways directions, so as to observe features when the vehicle banks. We assume perfect feature extraction and data association in the simulation in order to concentrate on the path planning and estimation results. The noise strength added to the simulated sensor data is shown in Table I.
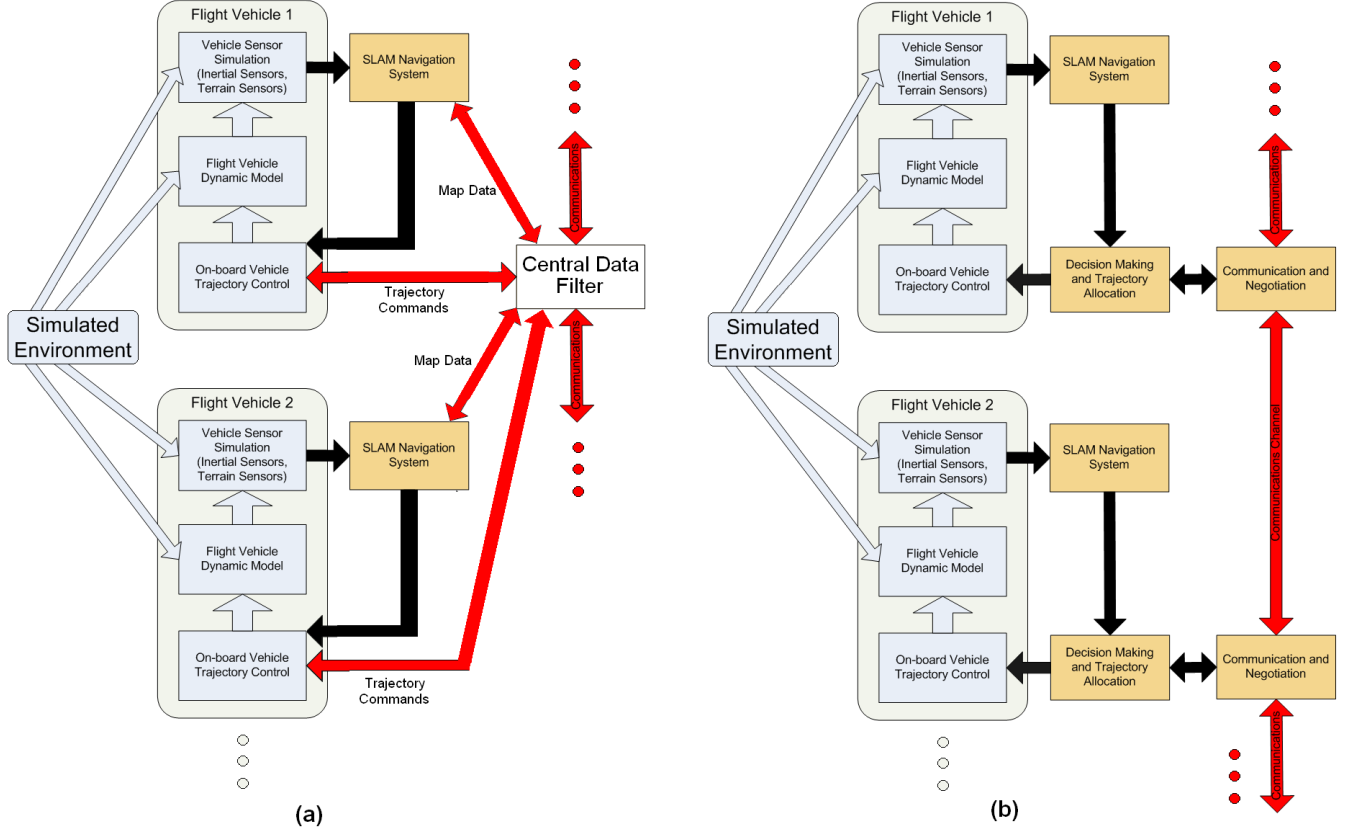
Fig. 2. Multi-UAV Co-operative SLAM Simulator Components: The simulator consists of a 6-DoF dynamic model of the UAV along with guidance and autopilot systems for low-level trajectory control. Vehicle localisation estimates are provided by the SLAM algorithm using simulated IMU and laser-vision system readings and are used for feedback for the low-level control. (a) Centralised Data Fusion and Control: Map information and utilities/trajectory commands are communicated between each vehicle via the central data filter at regular intervals. (a) Decentralised Data Fusion and Control: Map information and utilities/trajectory commands are communicated directly between each vehicle at regular intervals.

E. Single-Vehicle SLAM

Using the simulated IMU readings and terrain observations for each vehicle, the SLAM algorithm described in [4] is implemented on each vehicle to estimate its position, velocity and attitude and build up a map of features that it observes. Data from the IMU is read into the EKF at 100Hz and the EKF update cycle is run when feature observations are received (feature sensor operates at 10Hz). The estimated position, velocity and attitude are all sent to the guidance and autopilot systems to act as feedback for vehicle control.

F. Communications, Multi-vehicle SLAM and Multi-Vehicle Control

Both centralised and decentralised architectures for control and data fusion are considered in the simulation. In the centralised case, all of the vehicles communicate to a single central ground station which stores the central data filter for the distributed SLAM and coordinates commands to each vehicle; in the decentralised case vehicles share map information and negotiate control strategies with each other directly. The communications occur at two different intervals: SLAM map information is communicated between nodes at one second intervals and utility processing and trajectory commands are communicated between nodes every 10 seconds. We assume that each vehicle accurately knows its initial position and attitude w.r.t each other.

III. Multi-UAV SLAM Algorithms and Information Flow

In the multi-vehicle SLAM problem, the estimated state becomes the position, velocity and attitude of multiple vehicles and the positions of point features in the environment. In this section we analyse a distributed architecture for performing the data fusion.

A. Centralised Architecture with Distributed SLAM

Rather than communicate all of the sensor information to a central data fusion source, as is often done in multi-vehicle SLAM, the filter processing can be distributed amongst each of the UAVs. We will consider a distributed architecture in
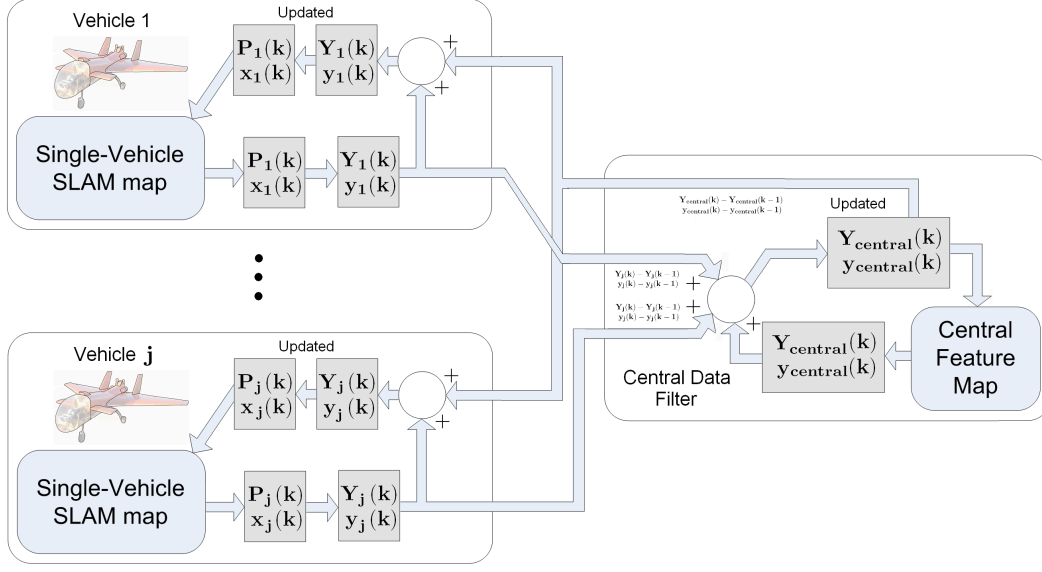
Fig. 3. Distributed Multi-vehicle SLAM: independent opinion pool architecture with local node feedback. Each vehicle communicates its posterior map estimates in information form which are added together at a central data filter. The central data filter then feeds back the information to each of the vehicles in order to update their local maps.

which each vehicle maintains its own local map and vehicle localisation estimates using an on-board implementation of the single-vehicle SLAM filter as shown in [4], but communicates the information relating to the feature map to a central source.

1) The Extended Information Filter and Inverse Covariance Form: The Extended Information Filter (EIF) [5] is a mathematically equivalent form of the EKF which uses the information matrix $\mathbf{Y}(k)$ and information vector $\mathbf{y}(k)$ to represent the estimate rather than the mean state vector $\mathbf{x}(k)$ and covariance matrix $\mathbf{P}(k)$. The relationship between the two forms is shown in Equations 1 and 2:

$$\mathbf{Y}(k) = \mathbf{P}^{-1}(k) \tag{1}$$
$$\mathbf{y}(k) = \mathbf{Y}(k)\mathbf{x}(k) \tag{2}$$

The advantage of the information or inverse covariance form of the EKF is that we can optimally combine two estimates of a state $(\mathbf{x}_1(k),\mathbf{P}_1(k))$ and $(\mathbf{x}_2(k),\mathbf{P}_2(k))$ together by simply adding their information matricies and adding their information vectors:

$$\mathbf{Y}_{combined}(k) = \mathbf{Y}_1(k) + \mathbf{Y}_2(k) \tag{3}$$
$$\mathbf{y}_{combined}(k) = \mathbf{y}_1(k) + \mathbf{y}_2(k) \tag{4}$$

provided that the errors in each estimate are not correlated to one another.

2) Distributed Data Fusion: The distributed data fusion is based on the independent opinion pool architecture shown in [6]. At regular intervals each vehicle takes its current state estimate relating to the map estimates only, i.e. $\mathbf{x}_m(k)$ and $\mathbf{P}_{mm}(k)$ where:

$$\mathbf{x}_m(k) = \begin{bmatrix} \mathbf{m}_1^n(k) \\ \mathbf{m}_2^n(k) \\ \vdots \\ \mathbf{m}_N^n(k) \end{bmatrix} \tag{5}$$

and $\mathbf{P}_{mm}(k)$ is a $3N \times 3N$ matrix of the elements of $\mathbf{P}(k)$ relating to the map feature estimates. Each vehicle then calculates its posterior information:

$$\mathbf{Y}_j(k) = \mathbf{P}_{mm}^{-1}(k) \tag{6}$$
$$\mathbf{y}_j(k) = \mathbf{Y}_j(k)\mathbf{x}_m(k) \tag{7}$$

for the jth vehicle where $j = 1, ..., M$ where M is the number of vehicles, and communicates this to the central map filter. The information that is sent will obviously be correlated to the information that was sent in the previous communication (since each vehicle's posterior information is based on the entire history of observations it has made). To overcome this, the central data filter maintains a record of the information that it has been sent in the previous communication $(\mathbf{Y}_j(k-1),\mathbf{y}_j(k-1))$ by each vehicle. When the new information arrives, the old information is subtracted from it before adding it to the central
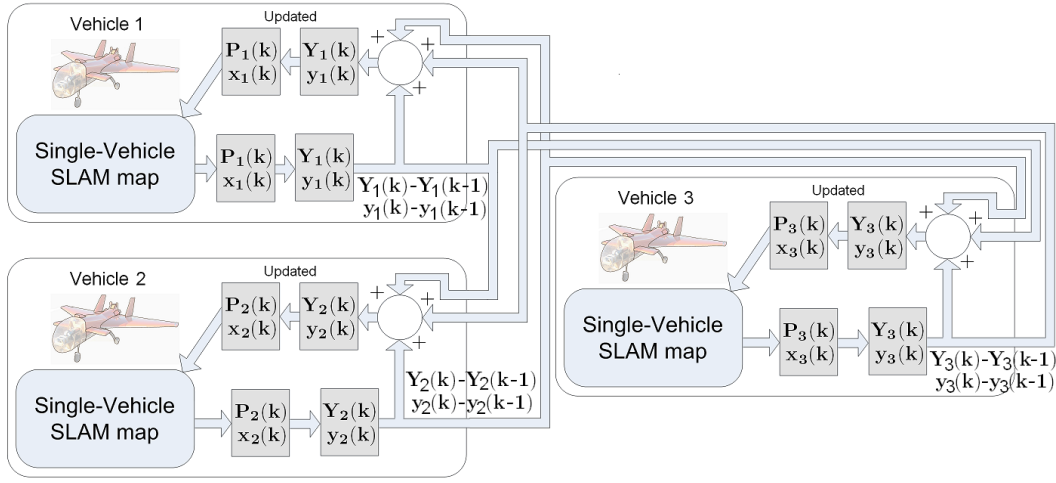
Fig. 4. Decentralised SLAM Data Fusion Architecture: Each UAV communicates to each other UAV in the team it's latest map state estimate in inverse covariance (information) form. Information is then added at the receiving end and converted from information space back into state space into the EKF.

map information, in order to remove the correlations and only count new information. The central map information update at the central data filter is thus:

$$
\mathbf{Y}_{central}(k) = \mathbf{Y}_{central}(k-1) + \ldots \tag{8}
$$
$$
\sum_{j=1}^{M}(\mathbf{Y}_j(k) - \mathbf{Y}_j(k-1))
$$
$$
\mathbf{y}_{central}(k) = \mathbf{y}_{central}(k-1) + \ldots \tag{9}
$$
$$
\sum_{j=1}^{M}(\mathbf{y}_j(k) - \mathbf{y}_j(k-1))
$$

Once the information is combined in the central filter, a state space estimate of the map feature locations and covariance can be recovered using Equations 10 and 11:

$$
\mathbf{P}_{mm,central}(k) = \mathbf{Y}^{-1}_{central}(k) \tag{10}
$$
$$
\mathbf{x}_{m,central}(k) = \mathbf{P}_{mm,central}(k)\mathbf{y}_{central}(k) \tag{11}
$$

3) Applying Local Node Feedback to the Independent Opinion Pool: So that each vehicle's localisation estimates can benefit from the observations of features made by other vehicles, information about the central map should be fedback to each of the local nodes. In the same way that was done on the central data filter, each vehicle must store the last information update that it received from the central filter ($\mathbf{Y}_{central}(k-1)$,$\mathbf{y}_{central}(k-1)$) so as not to double count the information that has been sent to it. Thus when each vehicle receives the communicated central information, its firstly computes its posterior information over the entire state space consisting of local vehicle estimate and map features using Equations 1 and 2 and updates this information using Equations 12 and 13:

$$
\mathbf{Y}_{local}(k) = \mathbf{Y}_{local}(k) + \ldots \tag{12}
$$
$$
(\mathbf{Y}_{central}(k) - \mathbf{Y}_{central}(k-1))
$$
$$
\mathbf{y}_{local}(k) = \mathbf{y}_{local}(k) + \ldots \tag{13}
$$
$$
(\mathbf{y}_{central}(k) - \mathbf{y}_{central}(k-1))
$$

The local information is then transformed back into state-space and covariance form to provide the updated estimate of the vehicle localisation and map features, which is substituted back into the EKF in the single-vehicle SLAM architecture. The operation of the central filter with local node feedback is illustrated in Figure 3.

This distributed architecture has several advantages over a completely centralised filter such as decreased required communication bandwidth (as only local estimates must be communicated, not observations and process model inputs) and the ability to deal with intermittent communications and delays as the information is maintained on the local vehicle.

B. Decentralised SLAM

The centralised filter shown above can be easily decentralised by removing the central data filter and having each local node perform the information fusion by receiving local information from each of the other vehicles.

At regular intervals each UAV takes its current state estimate relating to the map estimates and calculates its posterior information as shown in Equations 6 and 7. Each UAV maintains a record of the information sent during the last communication (i.e. $\mathbf{Y}_j(k-1)$,$\mathbf{y}_j(k-1)$) which is subtracted from the current information to form the new information that UAV has about the feature map [1]:

$$
\begin{aligned}
\mathbf{Y}_{j,new}(k) &= \mathbf{Y}_j(k) - \mathbf{Y}_j(k-1) & (14) \\
\mathbf{y}_{j,new}(k) &= \mathbf{y}_j(k) - \mathbf{y}_j(k-1) & (15)
\end{aligned}
$$

This new information is then communicated to each of the other UAVs. When each UAV receives all of the information updates from each of the other UAVs, this information is summed together along with the current UAV information to form the updated estimate of the map features in information form:

$$
\begin{aligned}
\mathbf{Y}_{j,update}(k) &= \mathbf{Y}_j(k) + \sum_{j=1}^{M} \mathbf{Y}_{j,new}(k) & (16) \\
\mathbf{y}_{j,update}(k) &= \mathbf{y}_j(k) + \sum_{j=1}^{M} \mathbf{y}_{j,new}(k) & (17)
\end{aligned}
$$

Once all of the information from other UAVs is combined in the update, a state space estimate of the map feature locations and covariance can be recovered back into the EKF using Equations 10 and 11. The operation of the decentralised SLAM filter is illustrated in Figure 4. Provided there are no communication delays or dropouts between the UAVs, the estimates in the decentralised form of SLAM are equivalent to the centralised form.

IV. MULTI-UAV ACTIVE SLAM CONTROL ARCHITECTURES

In this section we describe the decentralised trajectory planning algorithms for multiple UAVs.

A. Information Measures

Information measures are used to assess the accuracy in the SLAM estimates and can be used as utility functions for deciding control actions in active SLAM. We will consider to main measures; entropy and trace. The entropy $H(x)$ of a multivariate gaussian probability distribution can be calculated from its covariance matrix $\mathbf{P}$ or Fisher information matrix $\mathbf{Y}$ as follows:

$$
\begin{aligned}
H(x) &= \frac{1}{2}log[(2\pi e)^n |\mathbf{P}|] & (18) \\
&= \frac{1}{2}log[(2\pi e)^n \frac{1}{|\mathbf{Y}|}] & (19)
\end{aligned}
$$

Entropy is a measure of the compactness of a distribution and thus the informativeness. The entropy of a Gaussian distribution is proportionate to volume of the uncertainty ellipsoid made by the covariance matrix and thus also proportionate to the product of the covariance matrix eigenvalues.

Another measure which we will use is the trace of the covariance matrix of the SLAM estimates. The trace is equal to the sum of the eigenvalues of this matrix and is thus proportionate to the sum of the lengths of each axis of the uncertainty ellipsoid.

We can use entropy or trace as a utility function for multi-UAV trajectory control based on the expected value of the covariance or information matrix relating to a particular set of vehicle actions. In the case of single UAV planning, the utility associated with a given trajectory that a UAV could follow is the entropy or trace of the map estimates at the end of the trajectory. In the case of multi-UAV planning, the utility associated with a set of potential trajectories, one for each UAV to follow, is the entropy or trace of the map estimates at the end of the trajectory, after all vehicles have shared and combined the information they have received from the observations made along each trajectory. The aim of the planning is to minimise the entropy or trace and thus maximise the accuracy of the estimates.

B. UAV Trajectory Planning: Associating Utility to Trajectories

At a fixed time interval, each of the UAVs performs the following steps:

---

[1]Old information is removed before communications as to prevent 'double-counting' of any information that was previously sent.
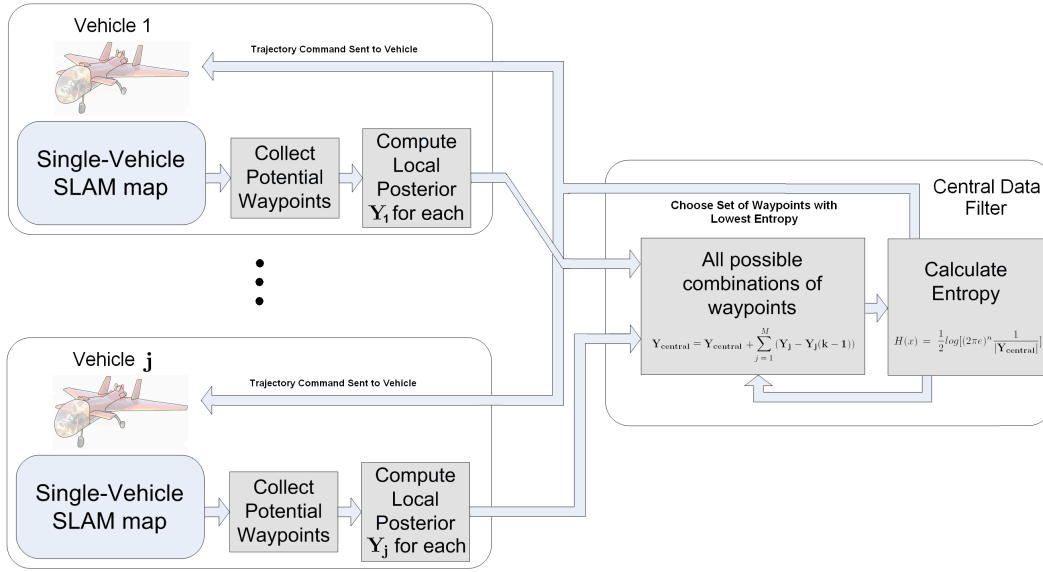
Fig. 5. Centralised Multi-UAV Active SLAM: At regular time intervals, each UAV evaluates several potential trajectories to take and the local posterior information that will be received for each. These plans are communicated to the central planning node. The central planning node then evaluates the updated information for each combination of trajectories, finds the optimal trajectory for each UAV and sends these trajectory commands to each vehicle.

1) Planning Potential Trajectories: The area around the vehicle out to 1km radius is broken up into 100×100m grids. At the center of each grid a potential waypoint is placed at a fixed altitude of 100m from the ground. The action space of potential trajectories is therefore the set of trajectories that consists of a steady turn followed by straight and level flight to each of the waypoints.

2) Approximating the Observations made along a Trajectory: Each UAV operates with three body-fixed, on-board terrain sensors: One sensor pointing downwards for observing features below the UAV during straight and level flight and two sideways pointing sensors; one pointing out to the right and the other out to the left of the UAV so as to observe features in the terrain below when the vehicle banks to turn. When solving the utility associated with a potential trajectory to take, the UAV makes an approximation of the observations it will make along the trajectory using the knowledge of the view-point constraints of each terrain sensor. Only features that have already been observed in the past are considered and new features that the vehicle may encounter during the trajectory are not taken into account.

3) Approximating the Posterior Information Matrix associated with a Trajectory: Using the expected trajectory data and approximated observations, an approximation of the posterior covariance at the end of the trajectory is computed by propagating the current time covariance along the trajectory segments using a discrete Ricatti equation [7]. Once we have the expected covariance, the expected local posterior information matrix is computed using Equation 6.

C. Centralised Multi-UAV SLAM

In the centralised control strategy, each UAV communicates both it's set of potential trajectories and the expected local posterior information at the end of the trajectory to the central control node. For each combination of trajectories, the central control node evaulates the utility of each and finds the set of trajectories which minimises the utility. The utility used in our case is the entropy of the expected posterior information matrix calculated by the central data node using Equations 8 and 19. The central command node then communicates the trajectory commands back to each of the vehicles to perform. The co-operative multi-UAV SLAM process is illustrated in Figure 5.

D. Co-ordinated Multi-UAV SLAM

In the co-ordinated control strategy, each vehicle individually computes the utility associated with each trajectory by substituting the local posterior information matrix for each potential trajectory into Equation 19. The vehicle then chooses the trajectory that minimises the utility (entropy). Since each UAV should have recently shared map information, each UAV should have the same map estimates and map estimate uncertainty before the planning stage. As each vehicle performs the trajectories it plans and shares the map information that it observes to other UAVs, this information is taken into account during the next planning time, and thus the UAV actions become co-ordinated in the task of minimising the shared map estimate uncertainty.
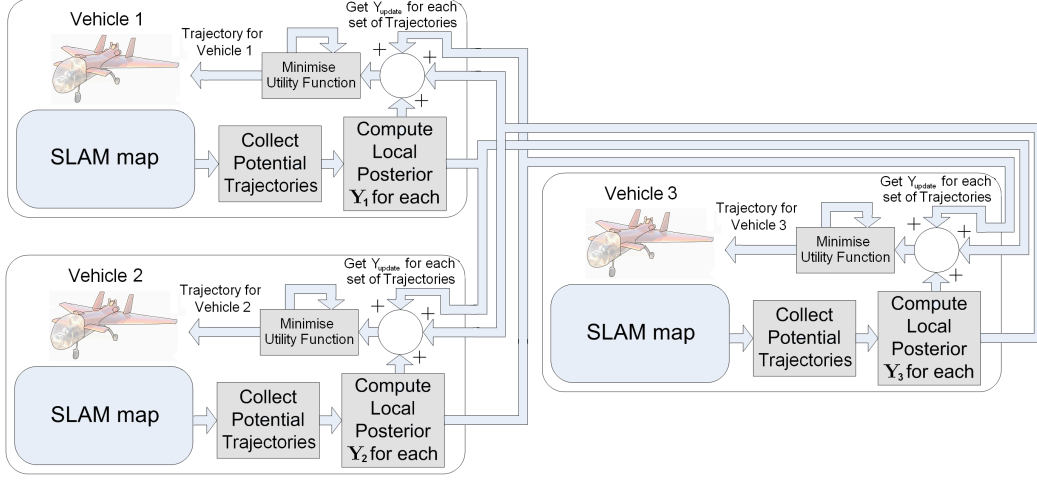
Fig. 6. Co-operative Multi-UAV SLAM: At regular time intervals, each UAV evaluates several potential trajectories to take and the local posterior information that will be received for each. These plans are communicated to each of the other UAVs. Each UAV then evaluates the updated information for each combination of trajectories and performs the trajectory that minimises the team map estimate uncertainty.

E. Co-operative Multi-UAV SLAM

In the co-operative control strategy, each UAV communicates both it's set of potential trajectories and the expected local posterior information at the end of the trajectory to each of the other UAVs. Each UAV now has all of the information necessary to determine the trajectory it should take that corresponds to the optimal team action. For each combination of trajectories from each of the other vehicles, each vehicle adds the expected local posteriors together (as shown in Equation 16) and computes the expected updated posterior information matrix for that potential team action. The utility of each potential team action is then calculated from Equation 19 and a team action with minimum entropy found. Each vehicle then takes its corresponding trajectory from the optimal team action it has calculated. It should be noted that given there are no communication losses between the vehicles the co-operative control actions will be identical to the centralised control actions (as shown in Section IV-C). The co-operative multi-UAV SLAM process is illustrated in Figure 6.

V. RESULTS

In this section we analyse simulation results concentrating on both the multi-vehicle SLAM algorithm and the coordinated and co-operative control architectures.

A. Multi-vehicle SLAM Results

In this subsection we concentrate on the results of multi-vehicle SLAM for three UAVs flying on fixed trajectories. Two simulation runs each for two different fixed trajectories were performed using the same features and sensor observations. A centralised data fusion architecture was used in the simulations. In the first run, none of the vehicles communicate with the central data filter and thus no map information is shared and each vehicle simply builds up a local SLAM of their surroundings. In the second run the vehicles now communicate with the central data filter at regular one second intervals in which map information is communicated.

Figure 7 illustrates the first set of trajectories taken by each vehicle in which each vehicle travels along the same large triangular path but where the starting point of each vehicle is offset from the others at each corner of the triangle. Figure 8 illustrates the second set of trajectories taken by each vehicle. Each vehicle follows a separate triangular segment path, where each of the triangles share two of their sides with each of the other triangles. For the second set of trajectories, each of the vehicles start in approximately the same location.

Figures 7 and 8 also show the resulting SLAM map and map feature estimate uncertainty ellipses at the end of the simulation. Values for the uncertainty ellipses are shown for both the communications and no communications cases. The ellipses have been shown at their 100σ level for clarity.

For the triangle segments trajectory (figure 8) we can see that each vehicle's local map is incomplete when there are no communications. This is due to because the coverage area of each vehicle does not overlap completely. The resulting map obtained when communications occur is complete (contains all of the features observed by each vehicle) due to map sharing. For the big triangle trajectory (figure 8) each vehicle follows the same path and thus eventually observes the whole map. The number of features observed in the communications and no communications scenarios are thus approximately the same for this set of trajectories.
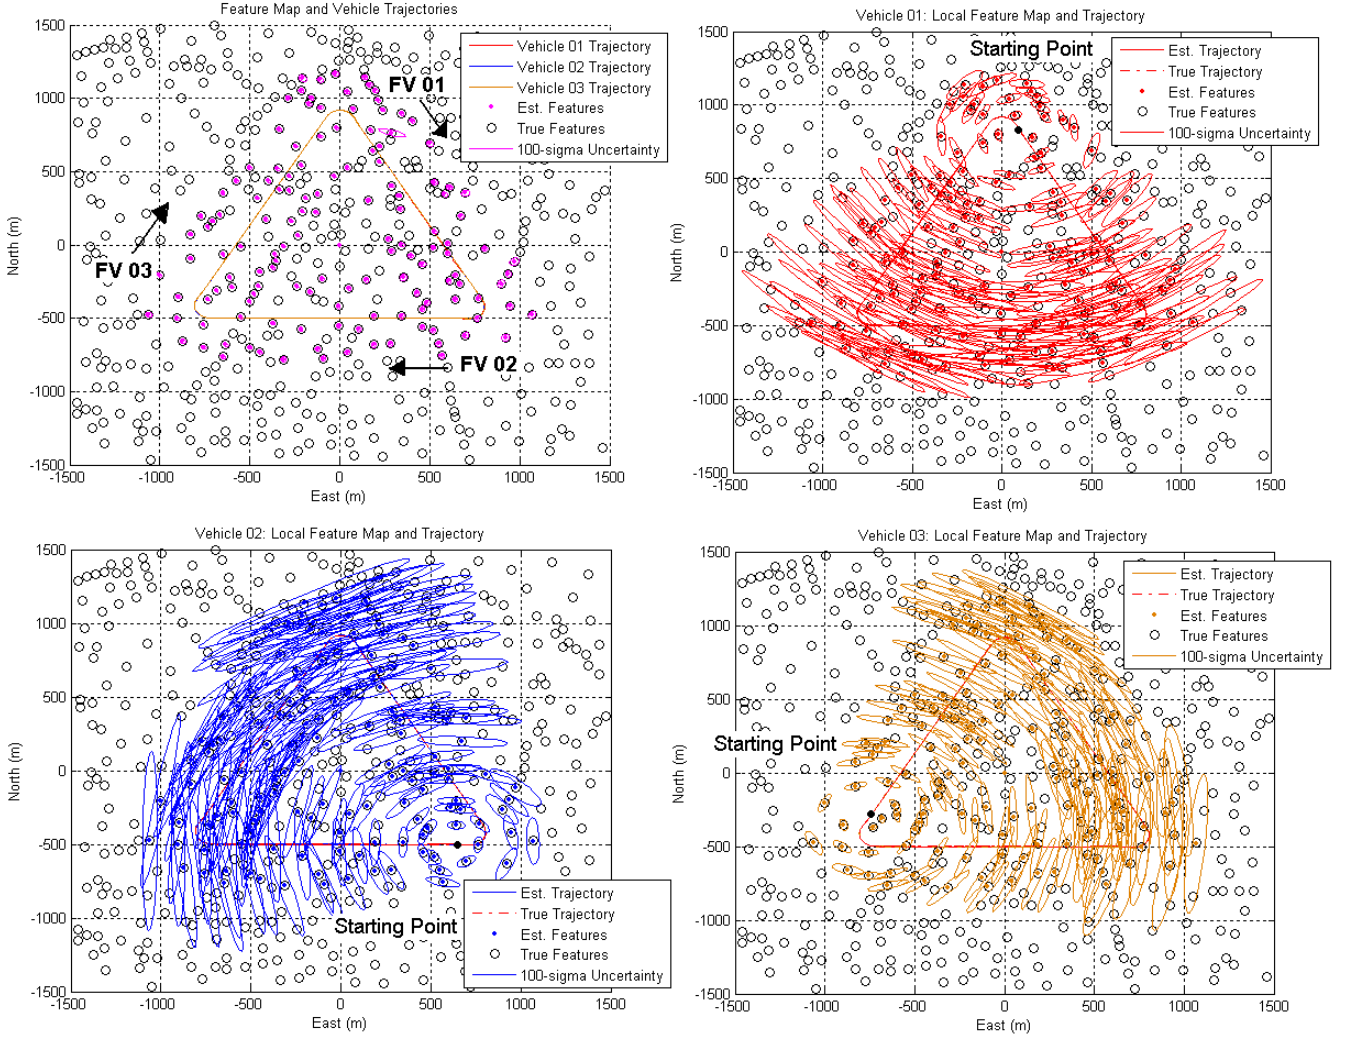
Fig. 7. Multi-UAV SLAM (Big Triangle Trajectory): Overhead view of the trajectories of each of the 3 vehicles and a comparison of the uncertainty ellipses for each of the three vehicle's map estimates when no communication is provided by the central data filter (blue, red and orange ellipses) and the uncertainty of the estimates when map information is shared amongst the vehicles via the central data filter (purple ellipses). The uncertainty ellipses are shown at their 100σ values. The uncertainty in the map estimates are greatly reduced when the vehicles share map information.

In both trajectory cases, the final map when vehicles have been communicating is significantly more accurate due to the combination of information from observations from multiple platforms. This accuracy is also increased due to the increase in localisation accuracy on the vehicle which has occurred due to the increase in available map information. In other words the final map accuracy for when information is continuously shared between vehicles is larger than the accuracy of just taking all of the information from each vehicle and the end of the trajectory in the case that vehicles don't communicate during flight. We can see that the uncertainty of each feature is consistently lower when the map information is shared amongst multiple vehicles via the central data filter.

Figures 9 and 10 illustrate the 3σ EKF covariance values for the vehicle position estimates from each of the three vehicles. A comparison is made between the uncertainty when no communications occur and when full communications of map information between the vehicles occurs via the central data filter. We can see that the uncertainty is decreased by the sharing of map information. This occurs due to the close coupling between localisation and mapping estimate accuracy in SLAM. When moving over features that are highly certain, the drift in the localisation system in well constrained and vice-versa when the localisation estimate accuracy is high, feature positions are initialised with a high degree of accuracy.

B. Multi-vehicle SLAM Results: Trajectory Comparison

In this section we compare the accuracy and information contained in the SLAM estimates for each of the two set trajectories shown in Figures 7 and 8. In these figures we can see that in the case of no communications, map feature position uncertainties are generally lower for features that are closer to the starting location of the vehicle than for those features further away. This
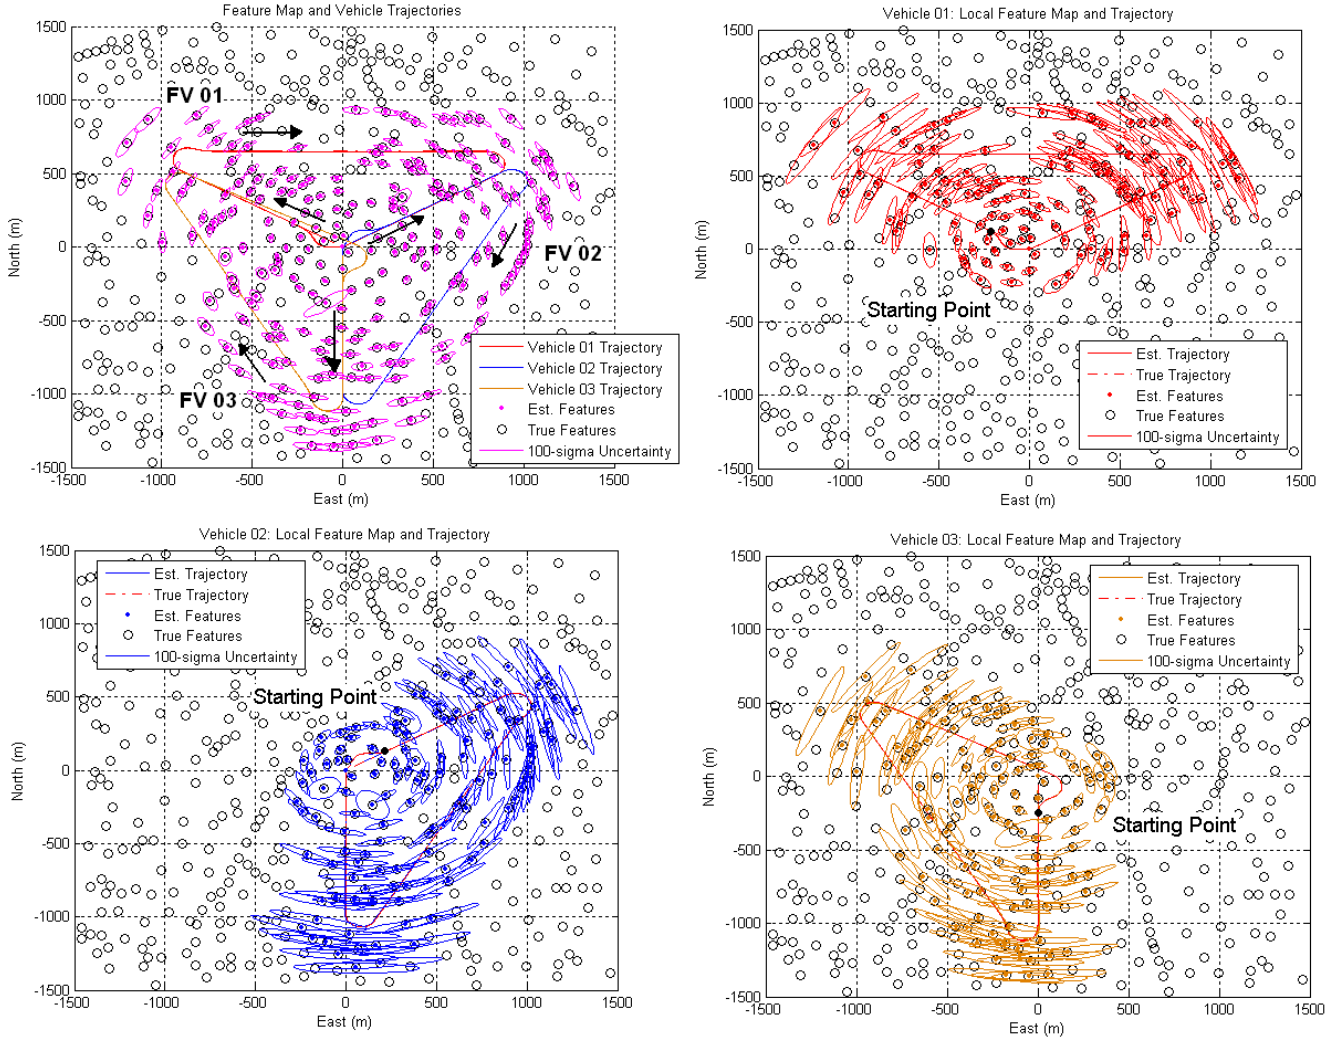
Fig. 8. Multi-UAV SLAM (Triangle Segments Trajectory): Overhead view of the trajectories of each of the 3 vehicles and a comparison of the uncertainty ellipses for each of the three vehicle's map estimates when no communication is provided by the central data filter (blue, red and orange ellipses) and the uncertainty of the estimates when map information is shared amongst the vehicles via the central data filter (purple ellipses). The uncertainty ellipses are shown at their 100σ values. The uncertainty in the map estimates are greatly reduced when the vehicles share map information.

is due to the build up in errors in localisation estimates that grow when the vehicle flies for a long time without revisiting terrain it has already mapped. When the vehicle completes a circuit and revisits features it has seen at the beginning of the trajectory, the uncertainty in all of the features in the map drops. This effect is known as 'closing the loop' where the vehicle location estimate is improved due to observing well known features. The uncertainty in other features in the map also drops due to the strong correlation with the vehicle location estimate. Subsequent visits to features far from the vehicle's starting location are required before their uncertainty can decrease to the level of features near the vehicle's starting location. The cycle of localisation uncertainty can be seen in Figures 9 and 10 as the vehicle follows around the loop several times through the entire flight.

Figures 11 and 12 illustrate the entropy, map size, map covariance matrix trace and trace per number of features for each of the two trajectories and for when communications and no communications occur. In Figure 11 and 12 we can see the sudden drop in covariance trace per number of features (equivalent to average feature uncertainty) around the 90 and 110 second mark for the case of no communication which corresponds to loop closing by each of the vehicles.

In both trajectory cases we can see that the map entropy and trace are much lower (thus higher map accuracy) for the communications case than for the case of no communications. Comparing the trace per number of features for each of the different trajectories we can see that the difference between the communications and no communications cases is much larger for the big triangle trajectory (figure 7) than the triangle segments trajectory (figure 8). This same trend can also be noticed in the size of the feature uncertainty ellipses for the communications cases in Figures 7 and 8 (purple ellipses). This is attributed to the fact that in the big triangle trajectory the vehicles all begin in different locations spread out over the map. When one
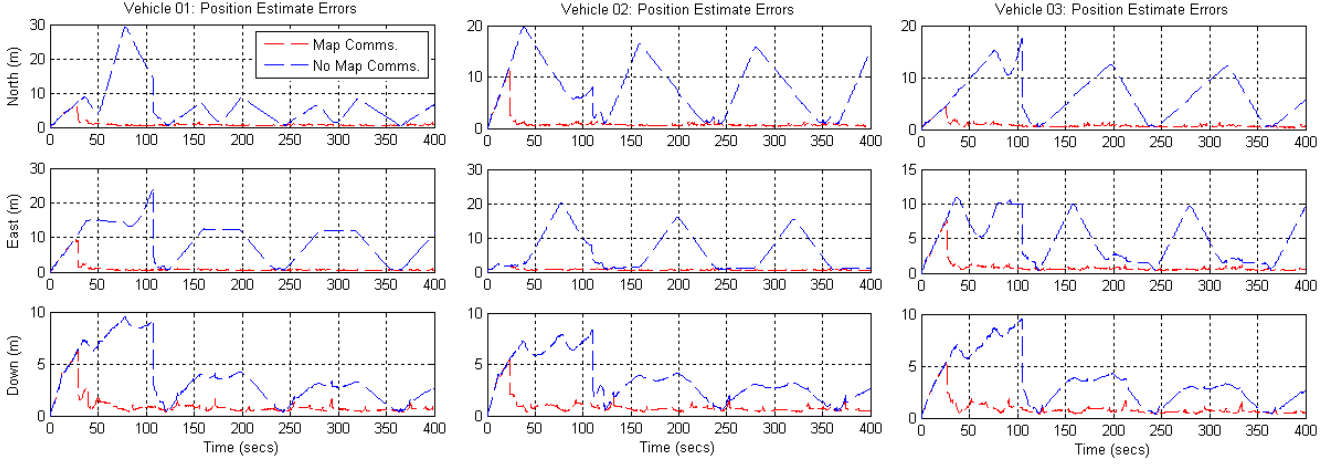
Fig. 9.   Multi-UAV SLAM (Big Triangle Trajectory): 3σ UAV position uncertainty. The solid red line indicates the uncertainty when the vehicle receives map updates from other vehicles observations. The dotted blue line indicates the uncertainty where no communications have taken place.
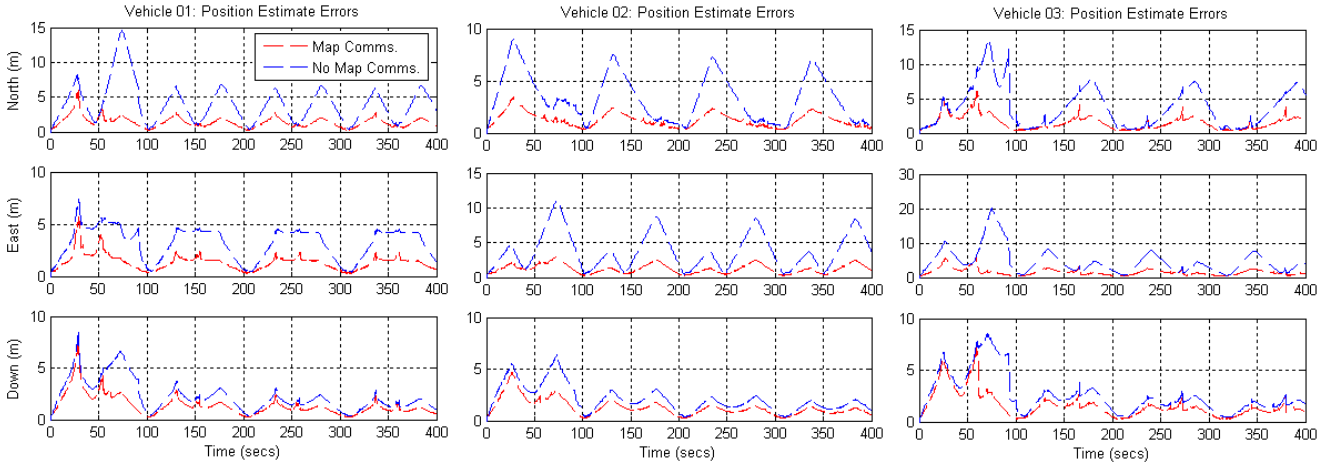

Fig. 10.   Multi-UAV SLAM (Triangle Segments Trajectory): 3σ UAV position uncertainty. The solid red line indicates the uncertainty when the vehicle receives map updates from other vehicles observations. The dotted blue line indicates the uncertainty where no communications have taken place.

vehicle flies over and observes the features at the starting location of one of the other vehicles, this helps to close the loop for this vehicle and thus reducing the uncertainty in the system dramatically. In the triangle segments trajectory case, each vehicle starts in the same local area and closes the loop over features that are further from the starting point of other vehicles. The strength of the loop closure and it's effect on reducing uncertainty is thus less. This can also be seen by the differences in the communications and no communications for the UAV position accuracies in Figures 9 and 10 where the position uncertainty for the communications case for the big triangle trajectory case is very small.

C. Co-operative vs. Co-ordinated SLAM: Confined Area Results

In this section we present results from the active SLAM algorithms presented in Section IV. In the following set of results the vehicle's flight trajectory was confined to a small sized area (approximately 800x800m) in order to compare more thoroughly the information gains in the co-operative and co-ordinated control architectures. Three different control strategies were compared. In the first test each vehicle performs SLAM and path planning independently of the other vehicles where no map information is shared across that network and each vehicle plans its trajectory based on maximising the information gain only in it's own local map. This control strategy is the same as that shown in the first year report [4] (i.e. single vehicle active SLAM). In the second test the vehicles take a co-ordinated approach to path planning where each vehicle shares map information across the network but path planning decisions are made locally without considering other vehicle's actions (see Section IV-D). In the third test the vehicles take a co-operative approach to path planning where map information is shared across the network and each vehicle plans across the network with other vehicles to arrive at the best team action (set of trajectories)(see Section IV-E).
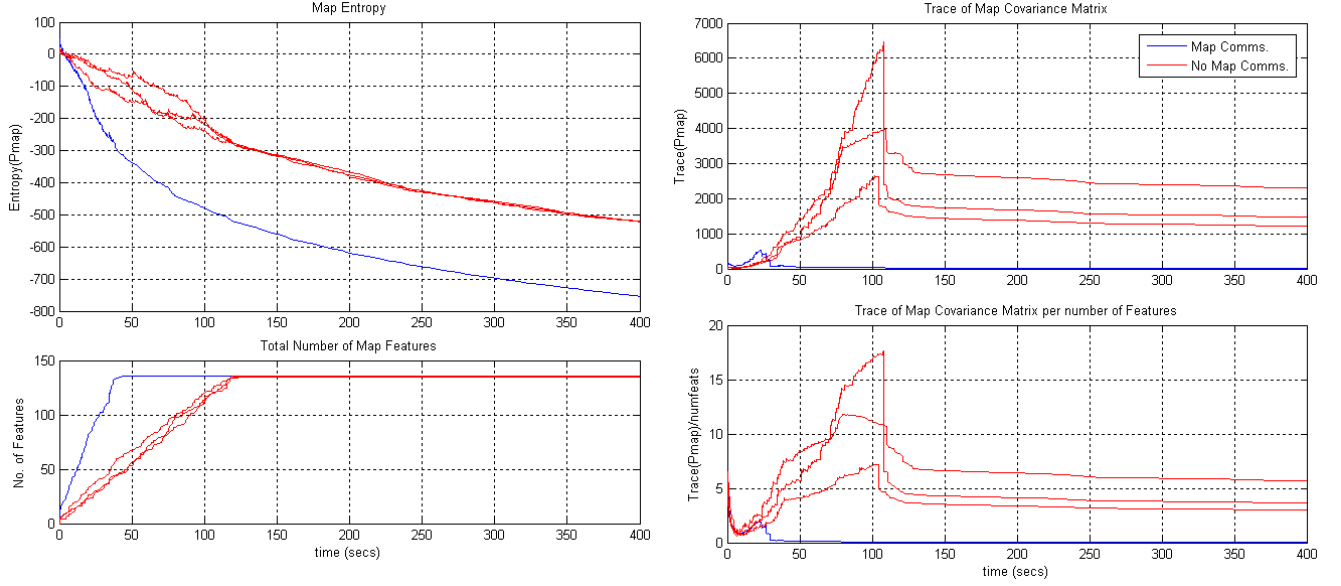
Fig. 11. Multi-UAV SLAM (Big Triangle Trajectory): Entropy, map size, covariance trace and covariance trace per feature number comparison for runs with map information communication and for runs with no map communication.
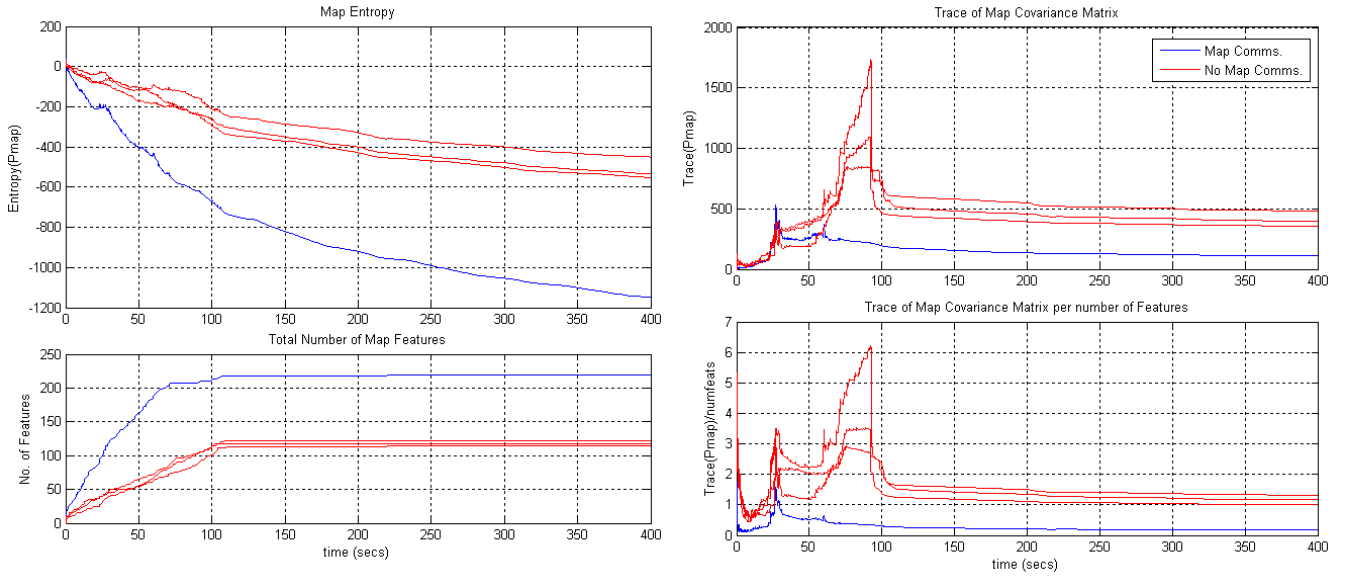

Fig. 12. Multi-UAV SLAM (Triangle Segments Trajectory): Entropy, map size, covariance trace and covariance trace per feature number comparison for runs with map information communication and for runs with no map communication.

1) Independent Vehicle SLAM and Path Planning Results: Figure 13 shows the trajectory taken by each of the three vehicles. Each of the vehicles tend to remain in their own local areas making observations of features in their own local maps. Figure 14 shows the resulting feature maps for each vehicle with associated 100σ uncertainty ellipses. As was seen in the results in section V-A, we can see that the uncertainty in the map estimates are large compared to those results shown in to next section due to the fact that map information has not been shared between the vehicles.

2) Co-ordinated Multi-UAV SLAM Results: Figure 15 shows the trajectories taken by each vehicle and the final map estimates for the co-ordinated control strategy where each vehicle makes trajectory plans at 5 second intervals. The trajectories of each of the vehicles have become more tightly integrated than in the independent planning case where each of the vehicles move to observe the same features. Figure 16 shows the trajectories taken by each vehicle and the final map estimates for the co-ordinated control strategy where each vehicle makes trajectory plans at 10 second intervals. When compared to the 5 second interval planning, the trajectories in this case follow a more consistent pattern. In the case of shorter planning times the vehicle tends to observe features that result in high information gain in the short term whereas with the longer planning timestep the
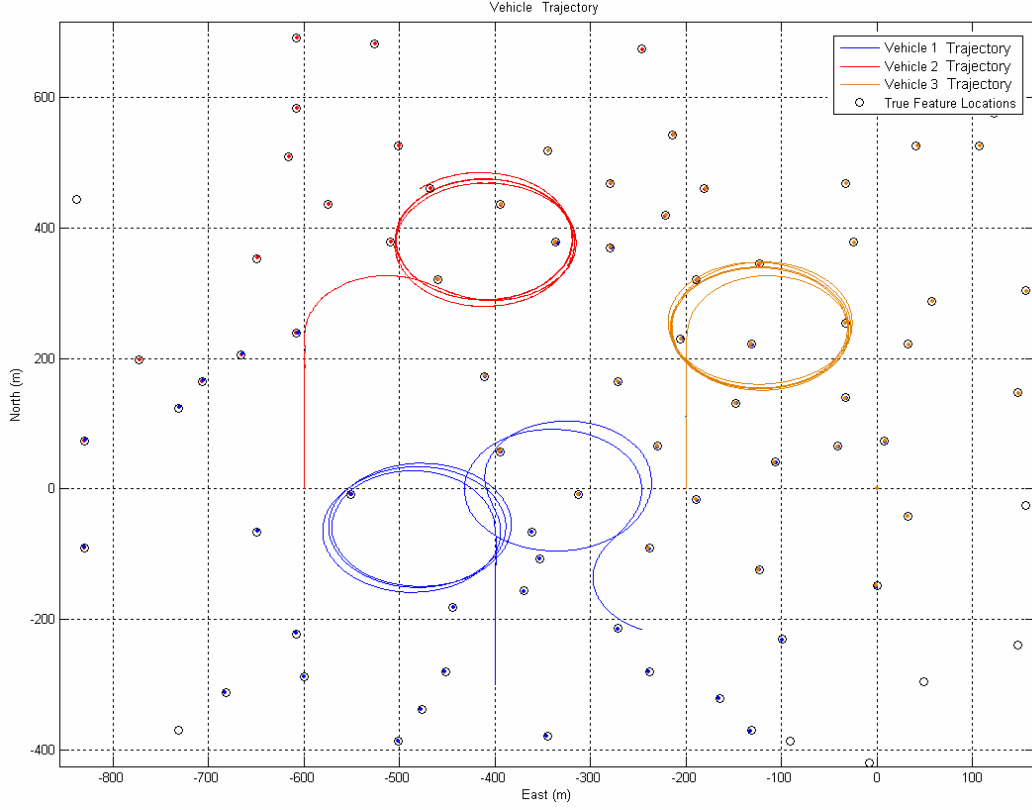
Fig. 13. Independent Vehicle SLAM and Path Planning: Shown are the trajectories of each of the vehicles where each vehicle does not share map information and trajectory planning is performed independently, without communicating to other vehicles.

trajectory is driven to maximise information gain over a longer time horizon where issues like closing the loop and localisation errors have more of an effect on the trajectory choice.

3) Co-operative Multi-UAV SLAM Results: Figures 17 and 18 show the trajectories taken by each vehicle and the final map estimates for the co-operative control strategy where each vehicle makes trajectory plans at 5 second and 10 second intervals. The trajectories are quite similar to the co-ordinated case, particularly for the 10 second time interval planning. The fact that the trajectories are so similar seems to indicate that over the confined area each vehicle's local optimal action is well matched to the team optimal action in terms of maximising map information.

Figures 19 and 20 show a comparison of the evolution of the entropy of the total map estimates for different control and data fusion architectures. We can see from Figure 19 that the entropy and thus uncertainty for the co-ordinated and co-operative cases is far lower than the uncertainty of the map estimates in the case of independent vehicle planning with no data fusion, thus highlighting the benefit of multi-UAV planning in SLAM. From Figure 20 we can see that the entropy is lower for the 10 second planning compared to the 5 second planning cases. The longer time step means that a more sensible trajectory is generated in the long term due to the increased level of fore-planning of the vehicle's actions. Figure 20 shows only marginal differences in the co-ordinated and co-operative strategies suggesting that over the small area considered the local optimal actions considered by each vehicle in the co-ordinated case matchup well with the team optimal actions considered by each vehicle in the co-operative case.

D. Co-ordinated Multi-UAV SLAM Results

Figure 21 shows the trajectory taken by each UAV over an 80 second flight using the co-ordinated control strategy. UAV 1 spends most of the flight in it's own area, building up the map accuracy but failing to realise that venturing over to the other segments of the map where UAVs 2 and 3 have visited will help to strengthen their estimates and in turn benefit the information in the shared map. Eventually after the map starts to be built and the information is shared between UAVs, the UAV actions begin to become more integrated and beneficial to the total map accuracy.

E. Co-operative Multi-UAV SLAM Results

Figure 22 shows the trajectory taken by each UAV over an 80 second flight using the co-operative control strategy. In this case the UAVs more commonly venture out into the areas already observed by other UAVs which helps to strengthen the
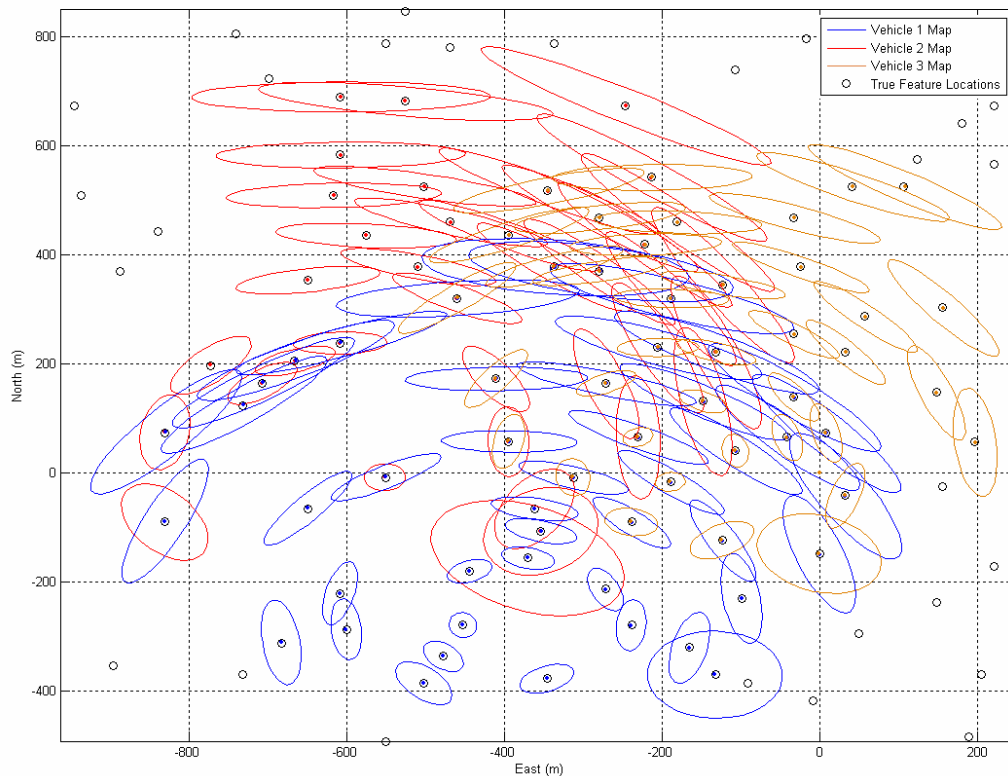
Fig. 14. Independent Vehicle SLAM and Path Planning: Shown are the map estimates and associated 100σ uncertainty ellipses for the local maps for each of the vehicles. Map information is not shared between vehicles and thus the uncertainty is higher than the results shown in Sections V-C.2 and V-C.3.

accuracy of the map due to the correlations that get built up between the map estimates from each of the UAVs.

VI. SUMMARY AND CONCLUSIONS

In this work we have examined the problem of co-ordinating the actions of multiple UAVs each performing inertial sensor based SLAM in order to maximise the accuracy of the shared terrain map the vehicles build and in turn increase the accuracy of each vehicle's own localisation states. In this report we have outlined algorithms for multi-UAV data fusion of SLAM map estimates using both centralised and decentralised architectures based on the extended information filter.

We examined three different UAV trajectory control architectures in order to maximise the information in the map estimates. The first control architecture was centralised. The second two architectures were decentralised where no central planner is required and each vehicle makes trajectory plans based on information that is shared between the UAVs. The first of the decentralised architecture used a co-ordinated approach where UAVs share map information in a data fusion sense but make their own locally optimal trajectory plans without considering the other vehicle's actions. The second decentralised architecture used a co-operative approach where vehicles negotiate over what trajectory each vehicle in the team should take in order to arrive at the team optimal action. In the case of large area operation, results showed that the co-operative approach seems to provide more sensible trajectories than the co-ordinated approach but at a higher cost in computation and communications bandwidth. Over a smaller sized area of operation, the co-ordinated and co-operative strategies seem to result in a similar set of trajectories being taken, with similar map information gain performance. This was attributed also to the fact that multi-UAV SLAM performed better when the vehicles were initially spread out from one another where each vehicle could contribute accurate estimates of the map around it's initial starting location.

REFERENCES

[1] M. Bryson, S. Sukkarieh, "Active Airborne Localisation and Exploration in Unknown Environments using Inertial SLAM," IEEE Aerospace Conference, Big Sky, 2006.

[2] J.H. Kim, S. Wishart, S. Sukkarieh, "Real-Time Navigation, Guidance and Control of a UAV using Low-cost Sensors," 4th International Conference on Field and Service Robotics, Yamanashi, 2003.

[3] D.T. Cole, S. Sukkarieh, A.H. Goktogan, H. Stone, R. Hardwick-Jones, "The Development of a Real-Time Modular Architecture for the Control of UAV Teams," 5th International Conference on Field and Service Robotics, Port Douglas, 2005.

[4] S. Sukkarieh, M. Bryson, R. Hardwick-Jones, "Cooperative Airborne Inertial-SLAM for Improved Platform and Feature/Target Localisation," Mid-term report, USAF/AOARD contract FA5209-05-P-0550, 2006.
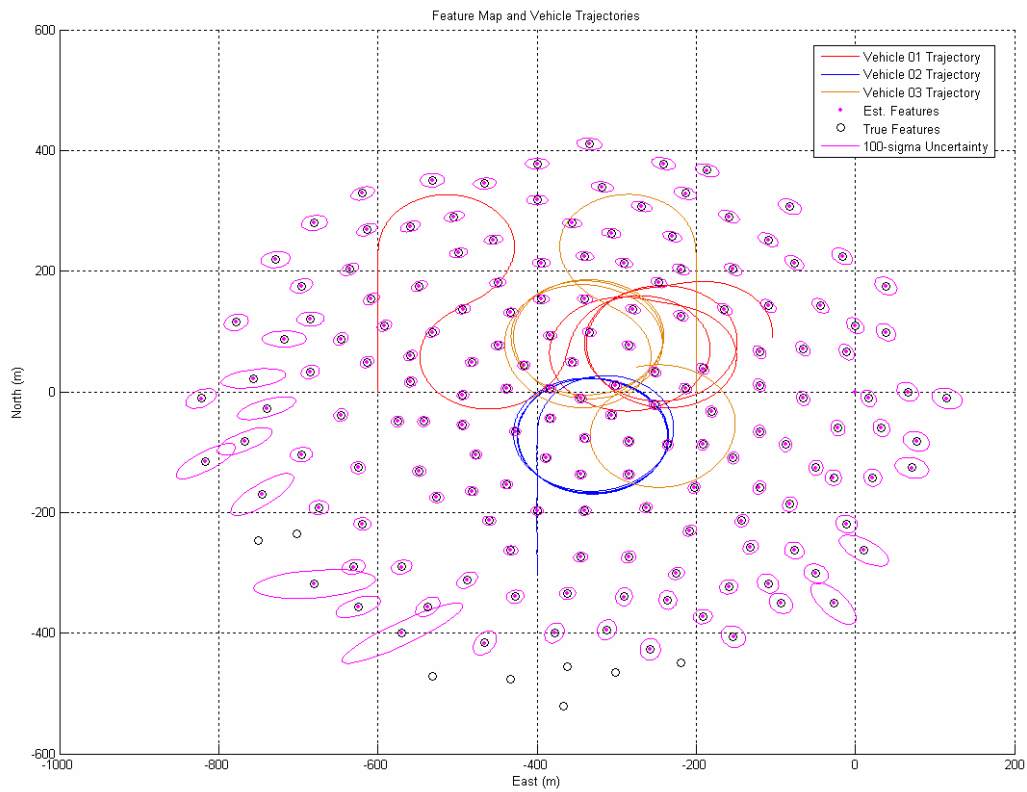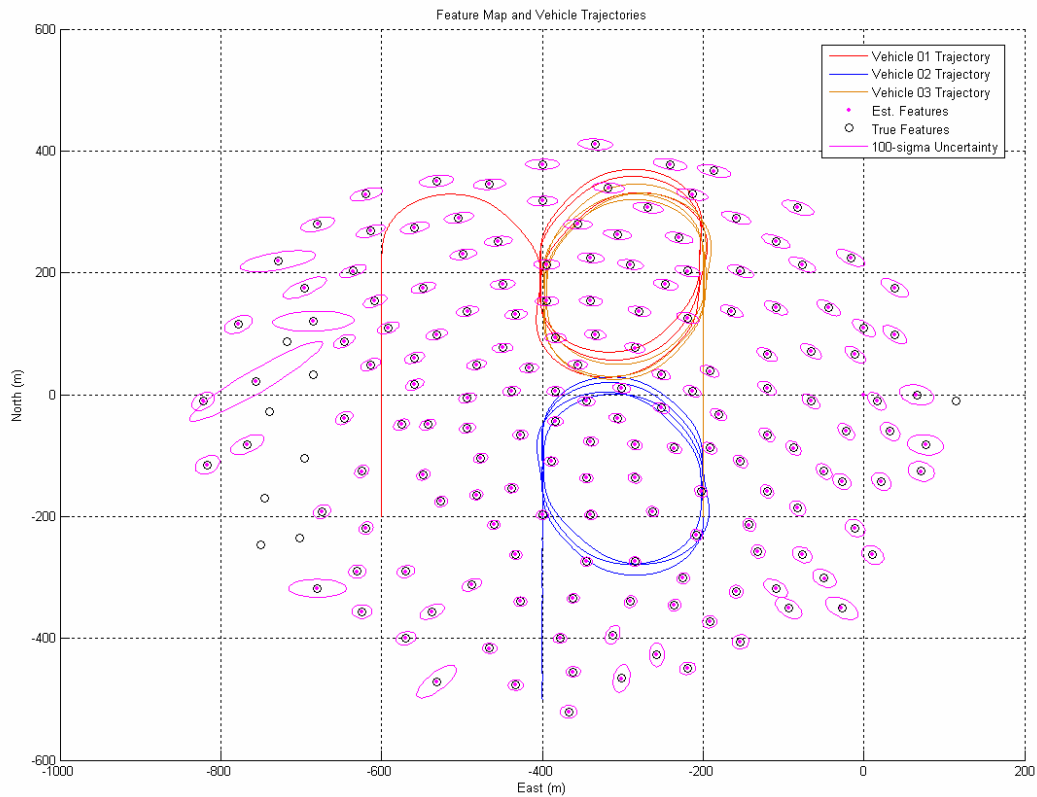
Fig. 15.   Co-ordinated Path Planning (5 second timestep): Shown is the trajectory taken by each vehicle and the 100σ uncertainty ellipses for the shared feature map.
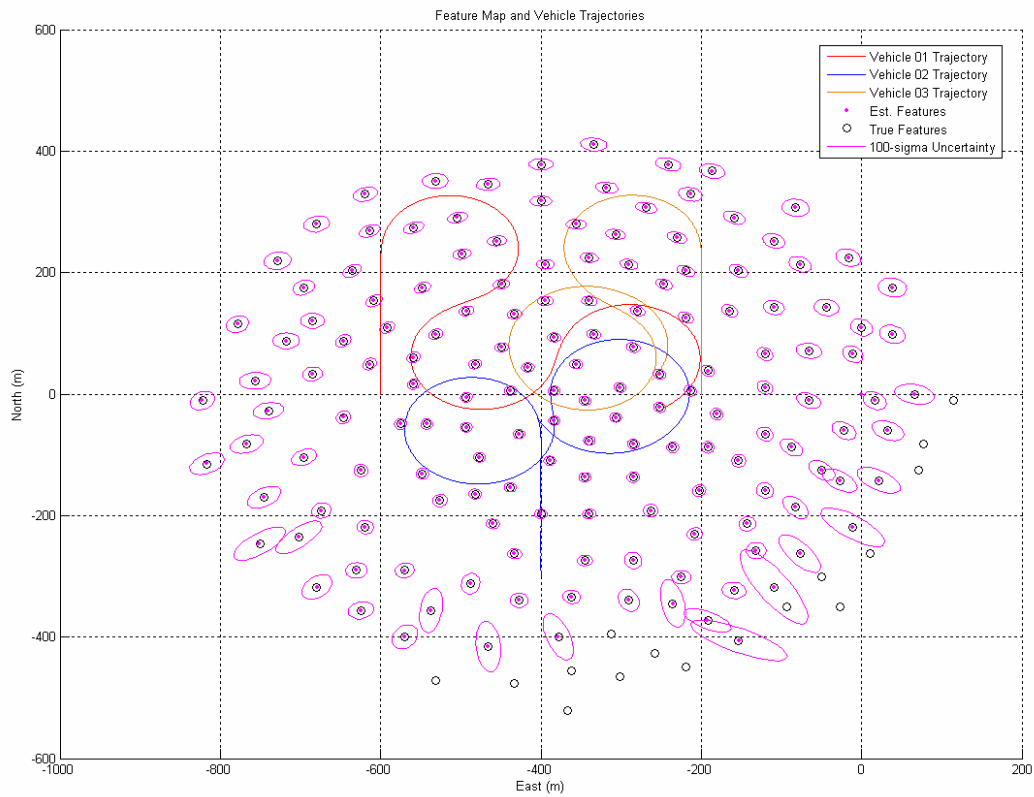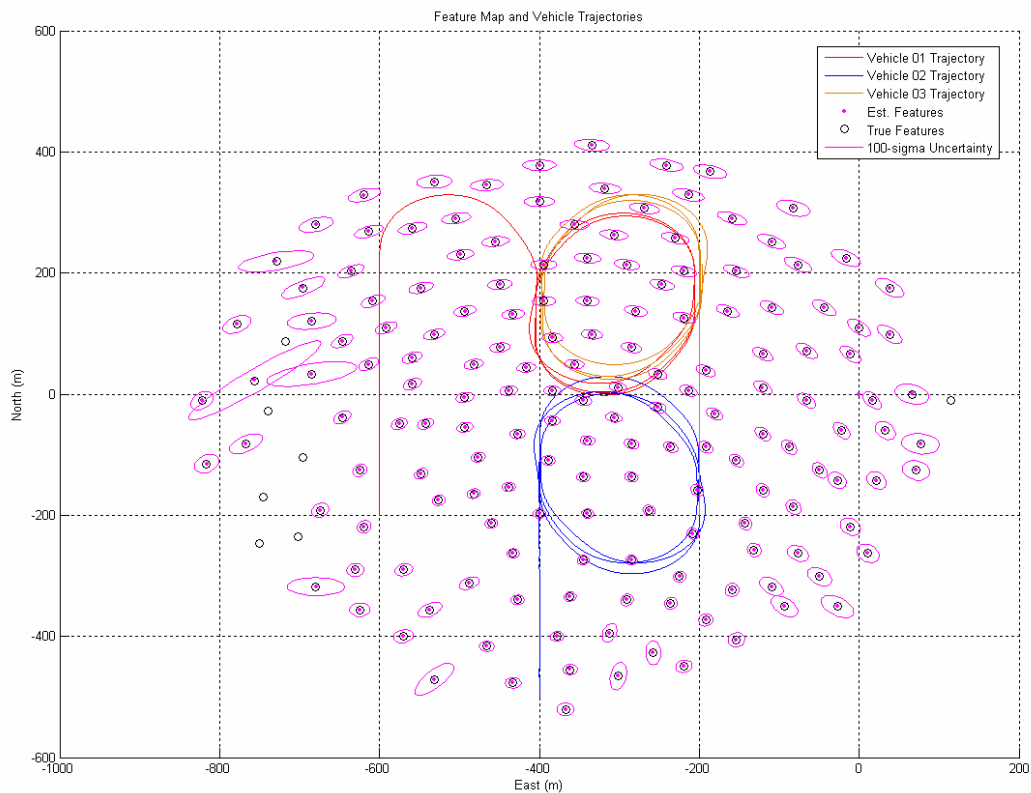
[5] P.S. Maybeck, "Stochastic Models, Estimation and Control, Volume 1," Academic Press, 1979.
[6] J. Manyika, H. Durrant-Whyte, "Data Fusion and Sensor Management: a decentralised information-theoretic approach," Ellis Horwood series in electrical and electronic engineering, 1994.
[7] M. Bryson, S. Sukkarieh, "Observability Analysis and Active Control for Airborne SLAM," To appear in IEEE Trans. on Aerospace and Electronic Systems, Vol. 43, No. 4, 2007.

Fig. 16. Co-ordinated Path Planning (10 second timestep): Shown is the trajectory taken by each vehicle and the 100σ uncertainty ellipses for the shared feature map.
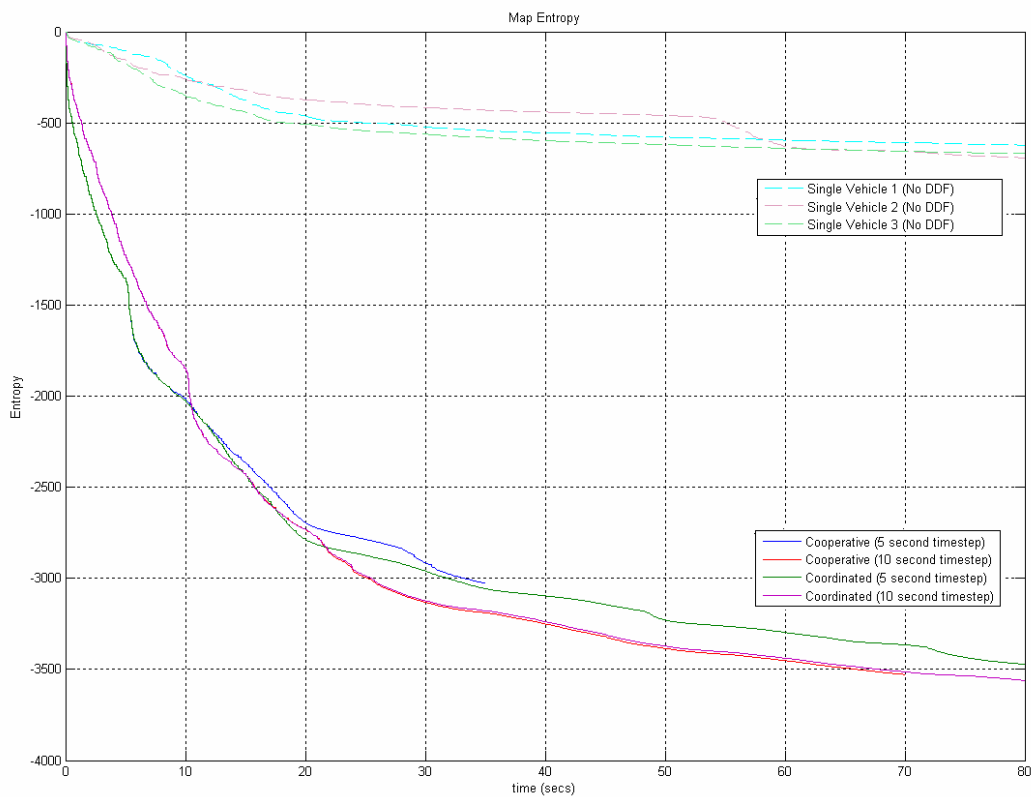

Fig. 17. Co-operative Path Planning (5 second timestep): Shown is the trajectory taken by each vehicle and the 100σ uncertainty ellipses for the shared feature map.

Fig. 18.   Co-operative Path Planning (10 second timestep): Shown is the trajectory taken by each vehicle and the 100σ uncertainty ellipses for the shared feature map.

Fig. 19.   Evolution of the entropy (uncertainty) in the map estimates for each of the different control architectures: The control architectures where map data sharing is present (co-ordinated and co-operative cases) have a much lower entropy than in the case of single vehicle independent planning where map data is not shared.
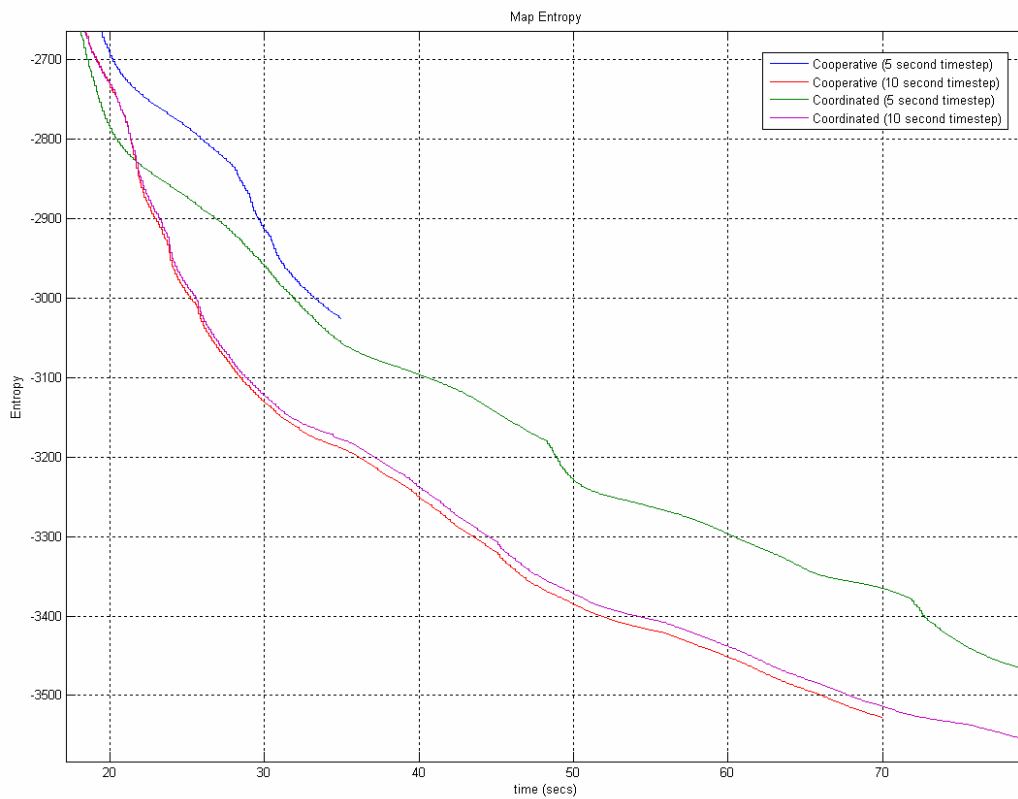
Fig. 20. Evolution of the entropy (uncertainty) in the map estimates for each of the different control architectures (zoomed in): the entropy difference between the co-operative and co-ordinated cases is small due to the small area of operation. Entropy is lower for the 10 second planning timestep cases due to the longer lookahead time of the planner.
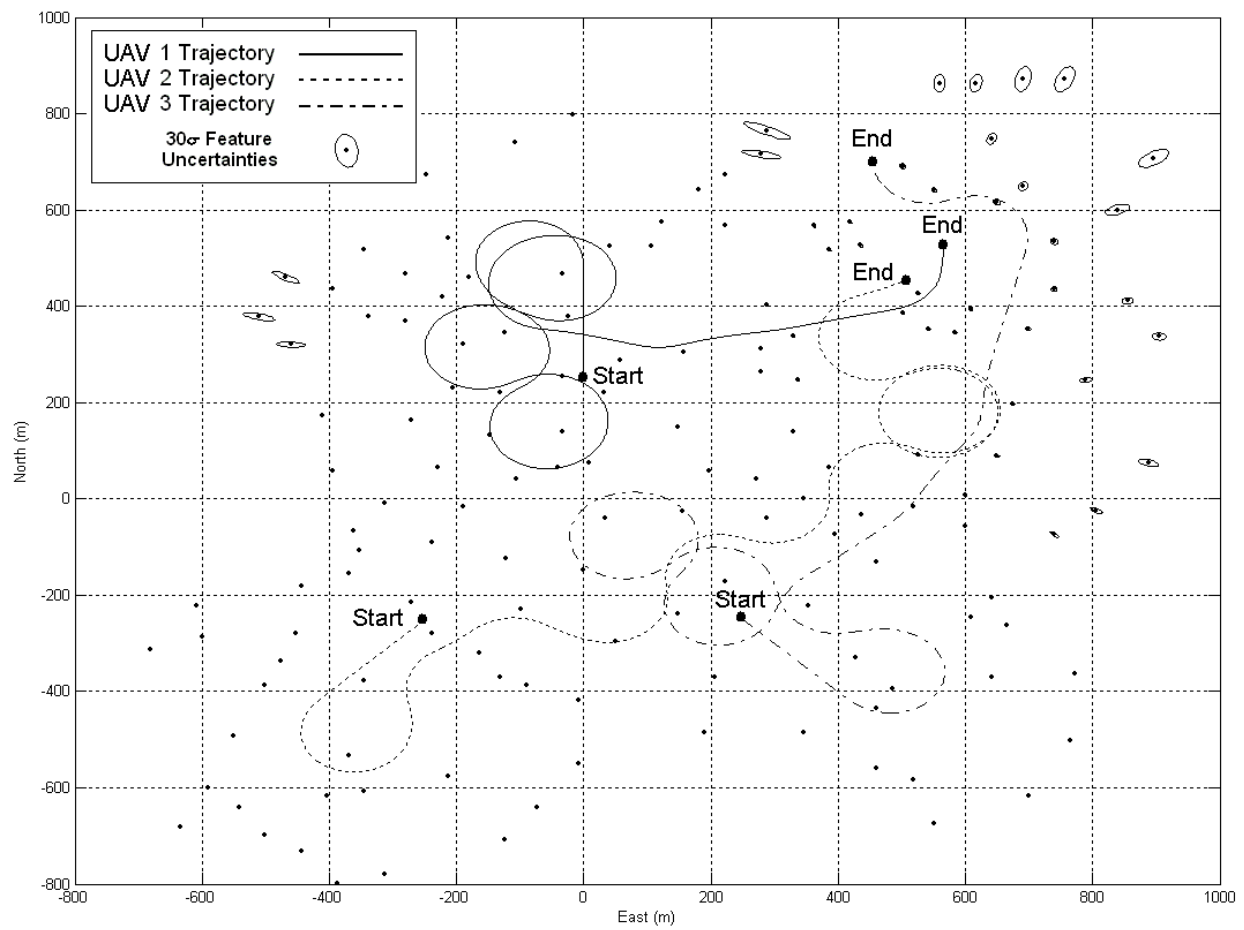
Fig. 21.   Co-ordinated Multi-UAV SLAM: Shown are the trajectories for each vehicle and the resulting map estimates at the end of the mapping process.
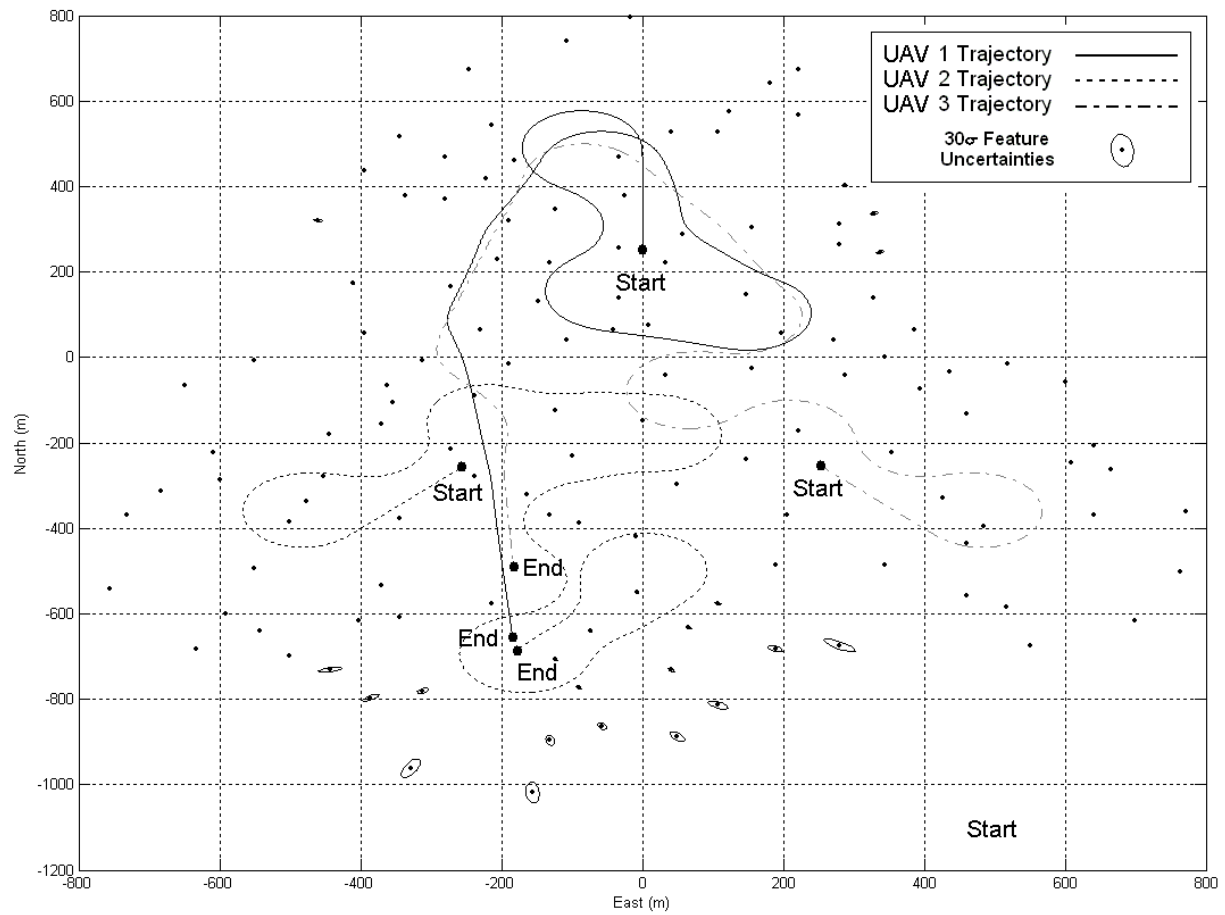
Fig. 22.   Co-operative Multi-UAV SLAM: Shown are the trajectories for each vehicle and the resulting map estimates at the end of the mapping process.